



OpenFabrics Alliance

Interoperability Working Group (OFA-IWG)

January 2010 Logo Event Report

UNH-IOL – 121 Technology Drive, Suite 2 – Durham, NH 03824 – +1-603-862-0090
OpenFabrics Interoperability Logo Group (OFILG) – ofalab@iol.unh.edu

Amit Krig
Mellanox Technologies
Hermon Building 4th Floor
P.O. Box 586, Yokonam 20692
Israel

May 11, 2010
Report Rev1.3
OFED Version: 1.5

Enclosed are the results from OFA Logo testing performed on the following devices under test (DUTs):

Mellanox Technologies MHGA28-XTC InfiniHost™ III Ex HCA card, dual-port 20Gb/s, PCIe1.2 x8 host bus, MemFree
Mellanox Technologies MHGS18-XTC InfiniHost™ III Lx HCA card, single-port 20Gb/s, PCIe1.2 x8 host bus, MemFree
Mellanox Technologies MHES14-XTC InfiniHost™ III Lx HCA card, single-port 10Gb/s, PCIe1.2 x4 host bus, MemFree
Mellanox Technologies MHRH29B-XTR ConnectX®-2 VPI card, Dual Port, 4X QSFP 20Gb/s InfiniBand
Mellanox Technologies MHZH29-XTR ConnectX®-2 VPI card, 4X QSFP 40Gb/s InfiniBand, SFP+ 10 Gigabit Ethernet
Mellanox Technologies MHQH29B-XTR ConnectX®-2 VPI card, Dual Port, 4X QSFP 40Gb/s InfiniBand
Mellanox Technologies MHGH29B-XTR ConnectX®-2 VPI card, Dual Port, 4X CX4 20Gb/s InfiniBand

The test suite referenced in this report is available at the IOL OFILG website, at test time release 1.30 (September 8, 2009 DRAFT) was used:

<http://www.iol.unh.edu/ofatestplan/>

Table 1 on the next page highlights the Mandatory test results required for the OpenFabrics Interoperability Logo for the DUT per the test plan referenced above and the current OpenFabrics Interoperability Logo Program (OFILP). Table 1 also lists “beta tests” which are not finalized and therefore not required for the logo.

For specific details regarding issues please see the corresponding test results.

Summary of all results follows on the second page of this report.

Testing Completed 02/09/2010

Nickolas Wood
ndv2@iol.unh.edu



Review Completed 05/11/2010

Bob Noseworthy
ren@iol.unh.edu

Table 1: Result Summary

The following table summarizes all results from the event pertinent to an IB device.

Note: **Failing a beta test does not necessarily indicate a failure of the device under test.** There are several reasons that a device might fail a beta test that are beyond the control of the device manufacturer. Such reasons include unresolved problems with the OFED software, with other vendors' devices, and with the test plan procedures.

Test Procedures	IWG Test Status	Result/Notes
10.1: IB Link Initialization	Mandatory	PASS
10.2: IB Fabric Initialization	Mandatory	PASS
10.3: IB IPoIB Connected Mode	Mandatory	PASS
10.4: IB IPoIB Datagram Mode	Mandatory	PASS
10.5: IB SM Failover and Handover	Mandatory	PASS
10.6: IB SRP	Mandatory	PASS
10.7: IB Ethernet Gateway	Beta	Not Tested
10.8: IB FibreChannel Gateway	Beta	Not Tested
12.1: TI iSER	Mandatory	Not Available
12.2: TI NFS over RDMA	Beta	Not Tested
12.10: HP MPI - HP	Beta	Not Tested
12.11: TI MPI - Intel	Beta	Not Tested
12.12: TI MPI - Open	Mandatory	PASS
12.13: TI MPI - OSU	Mandatory	PASS
12.3: TI Reliable Datagram Service	Mandatory	PASS
12.4: TI SDP	Mandatory	PASS
12.5: TI uDAPL	Mandatory	PASS
12.6: TI Basic RDMA Interoperability – Using Command Line	Mandatory	PASS with Comments
12.7: TI Basic RDMA Interoperability – Using XanStorm	Beta	Not Tested

Digital Signature Information

This document was created using an Adobe digital signature. A digital signature helps to ensure the authenticity of the document, but only in this digital format. For information on how to verify this document's integrity proceed to the following site:

http://www.iol.unh.edu/certifyDoc/certificates_and_fingerprints.php

If the document status still indicates “Validity of author NOT confirmed”, then please contact the UNH-IOL to confirm the document's authenticity. To further validate the certificate integrity, Adobe 6.0 should report the following fingerprint information:

MD5 Fingerprint: EEE1 7A82 7806 EB21 AF94 F189 E4BE 361B

SHA-1 Fingerprint: ECFB 7FAF AB4A 0832 2408 E965 9F5C E3F2 D784 AAAB

Report Revision History

- v1.0 Initial working copy
- v1.1 Updated to reflect retesting results. Added tested firmware md5sum.
- v1.2 Updated result tables
- v1.3 Updated topology

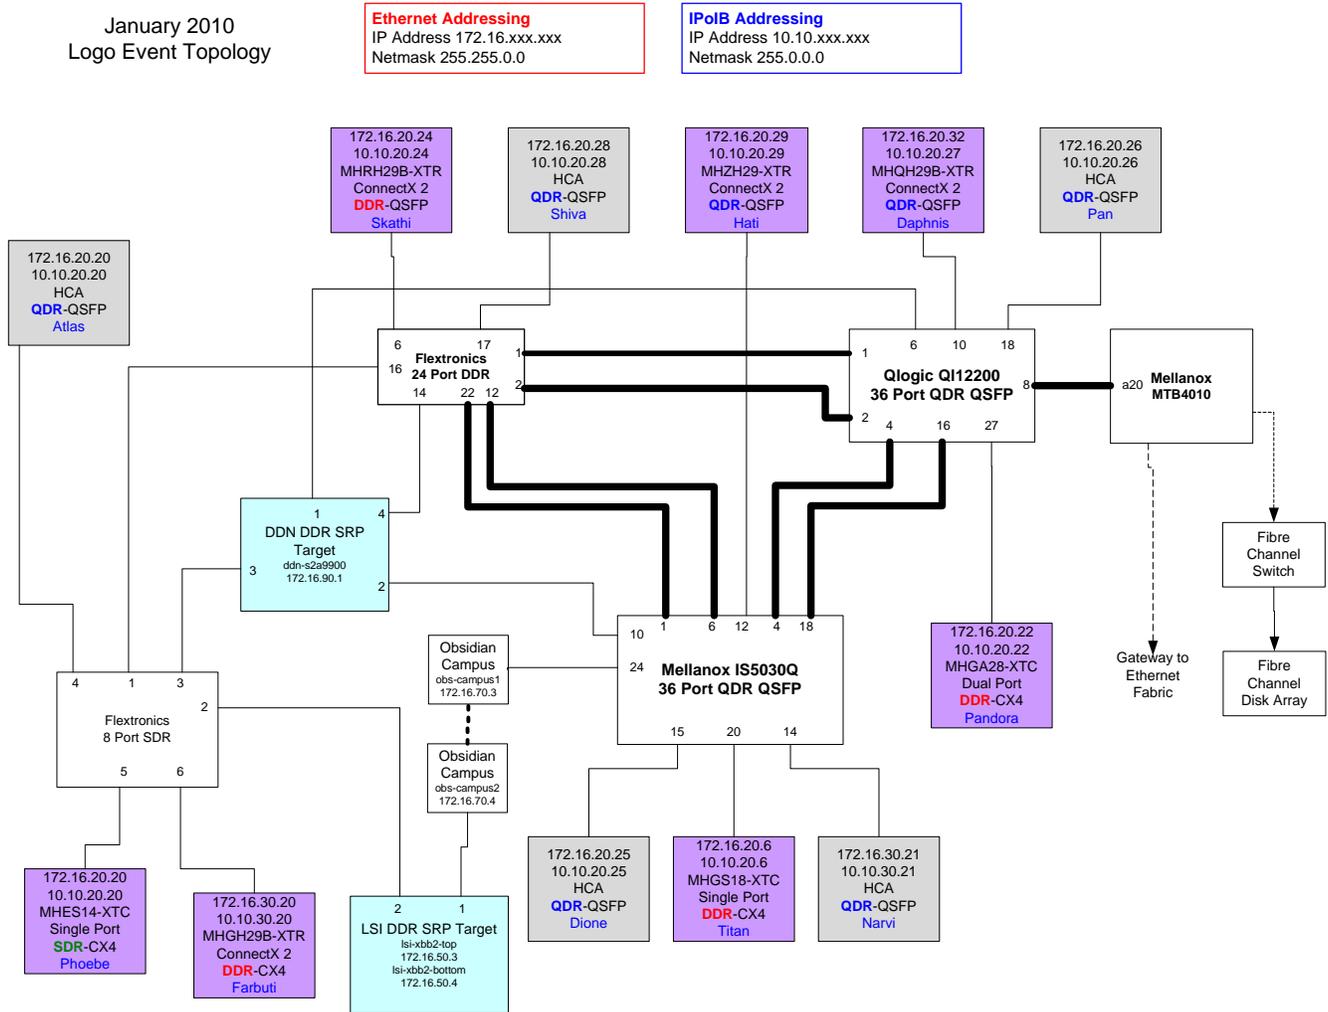
Table 2: Result Key

The following table contains possible results and their meanings:

Result:	Description:
PASS	The Device Under Test (DUT) was observed to exhibit conformant behavior.
PASS with Comments	The DUT was observed to exhibit conformant behavior however an additional explanation of the situation is included, such as due to time limitations only a portion of the testing was performed.
FAIL	The DUT was observed to exhibit non-conformant behavior.
Warning	The DUT was observed to exhibit behavior that is not recommended.
Informative	Results are for informative purposes only and are not judged on a pass or fail basis.
Refer to Comments	From the observations, a valid pass or fail could not be determined. An additional explanation of the situation is included.
Not Applicable	The DUT does not support the technology required to perform this test.
Not Available	Due to testing station limitations or time limitations, the tests could not be performed.
Borderline	The observed values of the specified parameters are valid at one extreme and invalid at the other.
Not Tested	Not tested due to the time constraints of the test period.

Table 3: DUT and Test Setup Information

Figure 1: The IB fabric configuration utilized for any tests requiring a multi-switch configuration is shown below.



DUT #1 Details			
Manufacturer	Mellanox	Firmware Rev	5.3.940
Model	MHGA28-XTC	Hardware Rev	20
Speed	DDR	Located in Host	Pandora
Additional Comments/Notes		Firmware bin md5sum	
MT25208		ccae1bef597d32ceebf19892b9eded4e	

DUT #2 Details			
Manufacturer	Mellanox	Firmware Rev	1.2.940
Model	MHGS18-XTC	Hardware Rev	20
Speed	DDR	Located in Host	Titan
Additional Comments/Notes		Firmware bin md5sum	
MT25204		200fa794e58245693fbd545d0bb4c9cc	

OFA Logo Event Report – January 2010
DUT: Mellanox HCAs

DUT #3 Details			
Manufacturer	Mellanox	Firmware Rev	1.2.940
Model	MHES14-XTC	Hardware Rev	a0
Speed	SDR	Located in Host	Phoebe
Additional Comments/Notes		Firmware bin md5sum	
MT25204		61ce4bc7e15892c7f1a22d95089342b9	

DUT #4 Details			
Manufacturer	Mellanox	Firmware Rev	2.7.200
Model	MHZH29-XTR	Hardware Rev	b0
Speed	QDR	Located in Host	Hati
Additional Comments/Notes		Firmware bin md5sum	
MT26428		4a9051ea457daea3465ebc9276b73cd9	

DUT #5 Details			
Manufacturer	Mellanox	Firmware Rev	2.7.200
Model	MHRH29B-XTR	Hardware Rev	b0
Speed	DDR	Located in Host	Skathi
Additional Comments/Notes		Firmware bin md5sum	
MT26418		fae90378f2c034724b733a5f407b7f99	

DUT #6 Details			
Manufacturer	Mellanox	Firmware Rev	2.7.200
Model	MHQH29B-XTR	Hardware Rev	b0
Speed	QDR	Located in Host	Daphnis
Additional Comments/Notes		Firmware bin md5sum	
MT26428		cbd0cd491629c217242b60f9cea13ce0	

DUT #7 Details			
Manufacturer	Mellanox	Firmware Rev	2.7.200
Model	MHGH29B-XTR	Hardware Rev	b0
Speed	DDR	Located in Host	Farbauti
Additional Comments/Notes		Firmware bin md5sum	
MT26418		c78322f0a1d83e4360524164a1601675	

Mandatory Tests - IB Device Test Summary Results:

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.1: IB Link Initialize	Test #1	No Issues seen	PASS
Discussion:			
Physical link initialization was verified between this device and every other device in the fabric. Link status was observed visually via status lights on the device. Port width and link speed was verified via ibdiagnet. Please see the table below for individual device results. A “P” indicates Passing, an “X” indicates that the Link Partner was the DUT itself and wasn’t tested since there was only 1 device of the model included, and a “RC” stands for Refer to Comments.			

Link Partner	DUT						
	MHGA28-XTC	MHGS18-XTC	MHES14-XTC	MHZH29-XTR	MHRH29b-XTR	MHQH29b-XTR	MHGH29b-XTR
Mellanox MHGA28-XTC DDR	X	P	P	P	P	P	P
Mellanox MHGS18-XTC DDR	P	X	P	P	P	P	P
Mellanox MHES14-XTC SDR	P	P	X	P	P	P	P
Mellanox MHZH29-XTR QDR	P	P	P	X	P	P	P
Mellanox MHRH29b-XTR DDR	P	P	P	P	X	P	P
Mellanox MHQH29b-XTR QDR	P	P	P	P	P	X	P
Mellanox MHGH29b-XTR DDR	P	P	P	P	P	P	X
Mellanox IS5030Q - QDR	P	P	P	P	P	P	P
QLogic 12200 - QDR	P	P	P	P	P	P	P
Flextronics X430044 - DDR	P	P	P	P	P	P	P
Flextronics X430066 - SDR	P	P	P	P	P	P	P
Obsidian Longbow	P	P	P	P	P	P	P
Obsidian Campus	P	P	P	P	P	P	P
LSI XBB2 - DDR	P	P	P	P	P	P	P
DDN S2A9900 - DDR	P	P	P	P	P	P	P
Mellanox MTB4010 - QDR	P	P	P	P	P	P	P

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.2: IB Fabric Initialization	Test #1:	No Issues seen	PASS
Discussion:			
The fabric configuration shown in Figure 1 was used for this test. 'ibdiagnet -c 1000' showed no Port error counters increment. Only one SM is run at a time. All switches are power cycled between SM trials. All links are validated via use of 'ibdiagnet' and 'ibchecknet' was used to verify that there were no duplicate guides.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.3: IPoIB Connected Mode	Test #1	No Errors Seen	PASS
	Test #2	No Issues Seen	PASS
	Test #3	No Issues Seen	PASS
Discussion: Test #1			
<p>Test #1: An automated test script was used to send ICMP Echo Request packets with payloads of specific sizes between all hosts on the configured fabric. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>Test #2: An HCA was disconnected from the fabric and reconnected in a different location; the ICMP Echo Reply packets ceased while the HCA was disconnected, and then resumed when it was reconnected. This procedure was repeated once with each subnet manager independently managing the fabric.</p> <p>Test #3: An automated test script was used to transfer a 4MB file using the SFTP protocol between all hosts on the configured fabric. The file was transferred once in each direction between all hosts, and the contents of the file were verified after each transfer. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>SMs tested include: <i>OFED OpenSM (SM Only)</i>, <i>QLogic 12200 (Managed Switch)</i>, and <i>Mellanox IS5030Q (Managed Switch)</i></p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.4: IPoIB Datagram Mode	Test #1	No Errors Seen	PASS
	Test #2	No Errors Seen	PASS
	Test #3	No Errors Seen	PASS
Discussion:			
<p>Test #1: An automated test script was used to send ICMP Echo Request packets with payloads of specific sizes between all hosts on the configured fabric. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>The last packet size tested; 65507 bytes, was seen to lose the very first packet sent. This can be fixed by setting the systems qlen value to 33 and as such has no bearing on these devices.</p> <p>Test #2: An HCA was disconnected from the fabric and reconnected in a different location; the ICMP Echo Reply packets ceased while the HCA was disconnected, and then resumed when it was reconnected. This procedure was repeated once with each subnet manager independently managing the fabric.</p> <p>Test #3: An automated test script was used to transfer a 4MB file using the SFTP protocol between all hosts on the configured fabric. The file was transferred once in each direction between all hosts, and the contents of the file were verified after each transfer. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>SMs tested include: <i>OFED OpenSM (SM Only)</i>, <i>QLogic 12200 (Managed Switch)</i>, and <i>Mellanox IS5030Q (Managed Switch)</i></p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.5: IB SM Failover and Handover	Test #1	No Issues Seen	PASS
Discussion:			
No issues were observed.			
SMs tested include: <i>OFED OpenSM (SM Only)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.6: IB SRP	Test #1:	No issues seen	PASS
Discussion:			
No SRP issues were observed with these HCAs			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.1: TI iSER	Test #1	Test Not Run	Not Tested
Discussion:			
This test was not performed due to the unavailability of an iSER target device in the fabric			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.12: TI MPI – Open	Test #1 PingPing and PingPong	No Issues Seen	PASS
	Test #2 All	No Errors Seen	PASS
Discussion:			
OpenMPI testing was only done in a Homogeneous fashion for this event. During the test a MPI ring was constructed using only HCA's Manufactured by Mellanox running on x86_64 systems.			
Notes			
Testing was done in two runs. Each run tested 5 HCA's by taxing the systems they reside in with 4 MPI processes each for a run total of 20 processes. Two runs were needed due to Mellanox providing 7 HCA's and our limitation of only being able to test 5 at a time.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.13: TI MPI – Mvapih1	Test #1 PingPing and PingPong	No Issues Seen	PASS
	Test #2 All	No Issues Seen	PASS
Discussion:			
MVAPICH testing was only done in a Homogeneous fashion for this event. During the test a MPI ring was constructed using only HCA Manufactured by Mellanox running on x86_64 systems.			
Notes			
Testing was done in two runs. Each run tested 5 HCA's by taxing the systems they reside in with 4 MPI processes each for a run total of 20 processes. Two runs were needed due to Mellanox providing 7 HCA's and our limitation of only being able to test 5 at a time.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.13: TI MPI – Mvapih2	Test #1 PingPing and PingPong	Test Not Run	Not Tested
	Test #2 All	Test Not Run	Not Tested
Discussion:			
This test was not run due to lack of Device Vendor support.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.3: TI Reliable Datagram Service	Test #1	No Issues Seen	PASS
	Test #2	No Issues Seen	PASS
Discussion:			
All tests succeeded from all DUT's without incident.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.4: TI SDP	Test #1: Netperf	No Issues Seen	PASS
	Test #2: SFTP	No Issues Seen	PASS
	Test #3: SCP	No Issues Seen	PASS
Discussion:			
No problems here observed with these DUT's.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.5: TI uDAPL	Test #1	No Issues Seen	PASS
Discussion:			
Every host was used as both a client and a server.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.6: TI Basic RDMA Interoperability – Using Command Line	Test #1	No Issues Seen	PASS with Comments
Discussion:			
Three HCAs needed an argument passed to the <code>ib_read_bw</code> command; specifically “-o 4”. These HCA’s; MHGA28-XTC (Pandora), MHGS18-XTC (Titan) and MHES14-XTC (Phoebe), are physically limited to a max of 4 outstanding requests.			
SMs tested include: <i>OFED OpenSM (SM Only)</i> , <i>QLogic 12200 (Managed Switch)</i> , and <i>Mellanox IS5030Q (Managed Switch)</i>			

Beta Tests - IB Device Test Results:

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.7: TI Basic RDMA Interoperability – Using XanStorm	Test #1	Not Tested	Not Tested
Discussion:			
Due to time constraints this test was not performed.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.10: MPI – Hewlett-Packard	Test #1 PingPing and PingPong	Test Not Run	Not Tested
	Test #2 All	Test Not Run	Not Tested
Discussion:			
This test was not run due to time constraints			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.11: TI MPI – Intel	Test #1 PingPing and PingPong	Test Not Run	Not Tested
	Test #2 All	Test Not Run	Not Tested
Discussion:			
This test was not run due to time constraints			